Low-rank approximations for large stationary covariance matrices, as used in the Bayesian and generalized-least-squares analysis of pulsar-timing data

Rutger van Haasteren and Michele Vallisneri
Jet Propulsion Laboratory, California Institute of Technology, Pasadena, California 91109
(Dated: July 11, 2014)

Many data-analysis problems involve large dense matrices that describe the covariance of stationary noise processes; the computational cost of inverting these matrices, or equivalently of solving linear systems that contain them, is often a practical limit for the analysis. We describe two general, practical, and accurate methods to approximate stationary covariance matrices as low-rank matrix products featuring carefully chosen spectral components. These methods can be used to greatly accelerate data-analysis methods in many contexts, such as the Bayesian and generalized-least-squares analysis of pulsar-timing residuals.

arXiv:1407.6710v1 [astro-ph.IM] 24 Jul 2014

I. INTRODUCTION

For various applications in the mathematical theory of Gaussian processes, in the field of machine learning, and (for these authors especially) in the data analysis of timing residuals from pulsars, one needs to take a matrix inverse of the form $(N + K)^{-1}$, where N is a very large diagonal matrix and K is a very large dense matrix obtained by evaluating a stationary correlation function k at a set of observation times t_i:

$$K_{ij} = k(t_i, t_j) = C(t_i - t_j) = C(\tau). \tag{1}$$

In the case of pulsar timing, the diagonal of the matrix N would contain the measurement errors for the times of arrival of the pulses, and $C(\tau)$ would describe correlated timing noise with a power-law spectrum $S(f)$, which is present in many pulsar-timing datasets [1]; by way of the Wiener–Khinchin theorem,

$$C(\tau) = \int_0^\infty \mathrm{d}f\, S(f) \cos(2\pi f \tau). \tag{2}$$

In the Bayesian analysis of pulsar-timing *residuals* [the times of arrival minus their theoretical model; see 2–4, and references therein] the inverse $(N+K)^{-1}$ is required to evaluate the likelihood of the data against a Gaussian-process model of its noise-like components [4]; in the generalized–least-squares analysis of residuals [5], the inverse enters the normal equations of the least-squares problem.

Taking the inverse of an $n \times n$ matrix is an $O(n^3)$ operation, which becomes forbidding for large n. A popular technique to speed up the inverse is to approximate the dense matrix K as a low-rank product $F\Phi F^T$, where F is $n \times m$ (with $m \ll n$) and Φ is diagonal (and $m \times m$). The Woodbury matrix identity [6] then yields

$$(N + K)^{-1} \simeq (N + F\Phi F^T)^{-1} = N^{-1} - N^{-1}F(\Phi^{-1} + F^T N^{-1} F)^{-1} F^T N^{-1}, \tag{3}$$

where the dominant cost is the $O(m^3)$ inversion of Φ. The generalized least-squares approach [5] uses a Cholesky decomposition [7] of $(N + K)$, which can also be obtained efficiently as a low-rank update [8–10]. The problem then is to obtain an adequate approximation in the form $F\Phi F^T$. A straightforward solution would be the truncated singular-value decomposition [7], which however is itself an $O(n^3)$ operation. Various other approximations have been discussed in the machine-learning literature [11]. This paper discusses the limitations of the popular "Fourier-sum" prescription used in pulsar timing (Sec. II), and presents two more accurate alternatives (Secs. III and IV).

In what follows, we characterize the accuracy of low-rank approximations by examining the difference between exact and approximated covariance matrices for power-law spectra of the form $S(f) = f^{-\gamma}\Theta(f - f_L)$, where $\gamma \in [1, 12]$ and where the low-frequency cutoff f_L is required to regularize the DC integral $C(0)$; for concreteness, we set $f_L = 1/(10\,T)$, where T is the total time span of observations. We also characterize the accuracy in approximating the *effective* covariance matrix that results after subtracting the best-fit constant, linear, and quadratic terms from the residuals. In pulsar-timing data analysis, this subtraction is performed to account for the intrinsic evolution of pulsar spins; the effective covariance matrix is given by PKP, where P is the orthogonal projector onto the null space of the *design matrix* given by constant, linear, and quadratic basis vectors evaluated at the observation times [12]. Along with this manuscript, we provide an implementation of these methods in the Python and C programming languages as ancillary files.

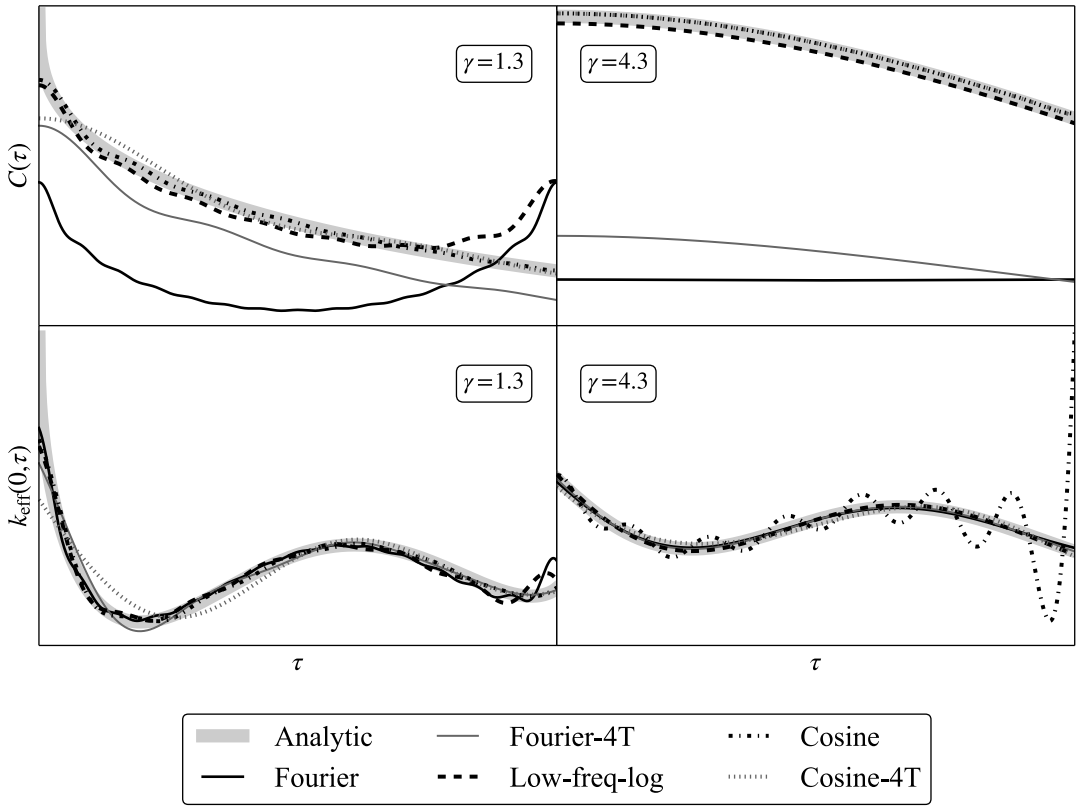

FIG. 1. Top: exact (thick gray) and approximated (all others, see legend) covariance function $C(\tau)$ for $S(f) = f^{-\gamma}\Theta[f - 1/(10T)]$ with $\gamma = 1.33$ (left) and $\gamma = 4.33$ (right). Bottom: exact and approximated *post-subtraction* correlation function $k_{\rm eft}(0,\tau)$ for the same models.

II. FOURIER-SUM EXPANSION OF COVARIANCE

In the context of pulsar timing, Lentati *et al.* [13] describe a method where correlated timing noise y is described as an explicit sum of Fourier components.

$$y(t) = \sum_{\mu=1}^{2q} w_\mu \phi_\mu(t) = \sum_{k=1}^{q} a_\mu \cos(2\pi k\, x) + b_\mu \sin(2\pi k\, x) \quad \text{with} \quad x = (t - t_0)/T, \tag{4}$$

where t_0 and T are the beginning time and duration of the observation. The w_μ are given a normal prior probability, leading to a diagonal matrix Φ where the same variance ρ_k is shared by the cosine and sine of the same frequency. The basis vectors $\phi_\mu(t_i)$ are only approximately orthogonal, because the t_i are not necessarily sampled regularly. In effect, Eq. (4) describes a Gaussian process with covariance

$$\begin{aligned} k(t_1, t_2) &= \sum_{\mu\nu} \phi_\mu(x_1) \Phi_{\mu\nu} \phi_\nu(x_2) \\ &= \sum_k \rho_k [\cos(2\pi k\, x_1)\cos(2\pi k\, x_2) + \sin(2\pi k\, x_1)\sin(2\pi k\, x_2)] \\ &= \sum_k \rho_k \cos[2\pi k (t_1 - t_2)/T] = \sum_k \rho_k \cos(2\pi f_k \tau) \end{aligned} \tag{5}$$

where $f_k = k/T$, for $k = 1, \ldots, q$. If we set $\rho_k = S(f_k)\Delta f = S(f_k)/T$, the sum produces a crude but practically useful approximation for the integral of Eq. (2). Thus we obtain a low-rank approximation $K \simeq F\Phi F^T$ where $F_{i\mu} = \phi_\mu(x_i)$ and $\Phi_{\mu\nu} = \delta_{\mu\nu}\rho_\mu$. Similar approximations have been studied in different contexts in the literature [14–16].

In Fig. 1 we show the exact and approximated $C(\tau)$ for $\gamma = 1.33$ and 4.33 (typical values in pulsar-timing applications), as well as the exact and approximated *post-projection* correlation function $k_{\rm eff}(0,\tau)$ [in this second case,

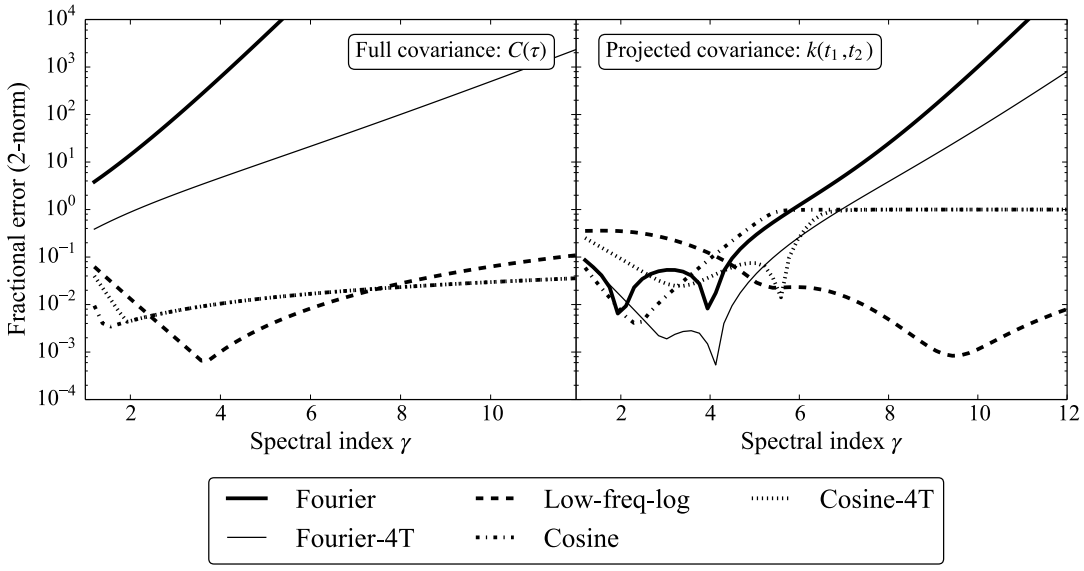

FIG. 2. Fractional approximation error as a function of γ for all approximation schemes discussed in this paper (see legend), before (left) and after (right) the subtraction of constant, linear, and quadratic trends. The fractional error is defined as the induced Euclidean norm (the 2-norm) of the difference between exact and approximated K_{ij} matrices, divided by the norm of the exact K_{ij}; it is evaluated for 1,000 equally spaced t_i between 0 and the maximum observation time T.

we cannot show $C(\tau)$ because the correlation function is not stationary after the subtraction of linear and quadratic trends]. The exact $C(\tau)$ and $k_{\rm eff}(0,\tau)$ appear as thick gray curves, while their Fourier-sum approximations are shown as thin dark continuous curves. The thin light continuous curves show the "$4T$" Fourier-sum approximation obtained by taking $f_k = k/(4T)$. The other curves correspond to the novel approximation schemes described in Secs. III and IV. For all of them we use seven frequency components, corresponding to F and Φ matrices with 14 columns. We are purposefully including fewer components that would be used in practice, so that the ways in which different approximations fail can be distinguished in the figure.

In Fig. 2 we display the fractional approximation error, defined as the induced Euclidean norm (the 2-norm) of the difference between exact and approximated K_{ij} matrices, divided by the norm of the exact K_{ij}; the fractional error is shown as a function of γ both before (left) and after (right) projection, and it is evaluated for 1,000 equally spaced t_i between 0 and the maximum observation time T. Again the thin dark and light continuous curves show the Fourier-sum approximation and its "$4T$" version. In this case, we use 30 frequency components for all approximation schemes, corresponding to F and Φ matrices with 60 columns. This many components appear to saturate the accuracy of the schemes, at least for γ larger than 3–4.

The top panels of Fig. 1 show that the Fourier-sum expansion is very inaccurate in approximating $C(\tau)$, because it lacks a constant offset term. In addition, the Fourier-sum expansion will always produce a *periodic* covariance function, although this is remedied in the "$4T$" version by pushing the periodic boundary way beyond the $[0, T]$ range over which the covariance function would be sampled. As shown in the bottom panels of Fig. 1, these handicaps disappear for the post-subtraction effective covariance, especially with the "$4T$" Fourier sum, because the subtraction removes the constant $C(\tau)$ offset and makes the approximation non-periodic. Figure 2 confirms this qualitative picture across the entire range of γ, and shows that for the projected covariance the "$4T$" Fourier sum is actually rather competitive among the approximations discussed in this paper.

III. LOW-FREQUENCY-LOGARITHMIC EXPANSION OF COVARIANCE

The top-left panel of Fig. 3 visualizes the approximation of Eq. (5) for the integral $C(0) = \int S(f)\,{\rm d}f$ as the rectangular sum $\sum_{k=1}^{q} S(f_k)\Delta f$, with $f_k = k/T$ and $\Delta f = 1/T$. As in the figure, we may associate the $S(k/T)$ samplings of the spectral density with the right edges of the rectangles; however, we may also center the rectangles, showing that the $[0, 1/2T]$ interval – just where red spectra blow up – is effectively left uncovered. Our improved approximation (first conceived by the first of these authors, RvH) consists in adding a few logarithmically spaced rectangles to this low-frequency interval, as shown in the right panel of Fig. 3, and in using a slightly more sophisticated integration rule. The bottom panels of Fig. 3 show the approximation of the oscillatory $C(\tau)$ integral, for a representative choice

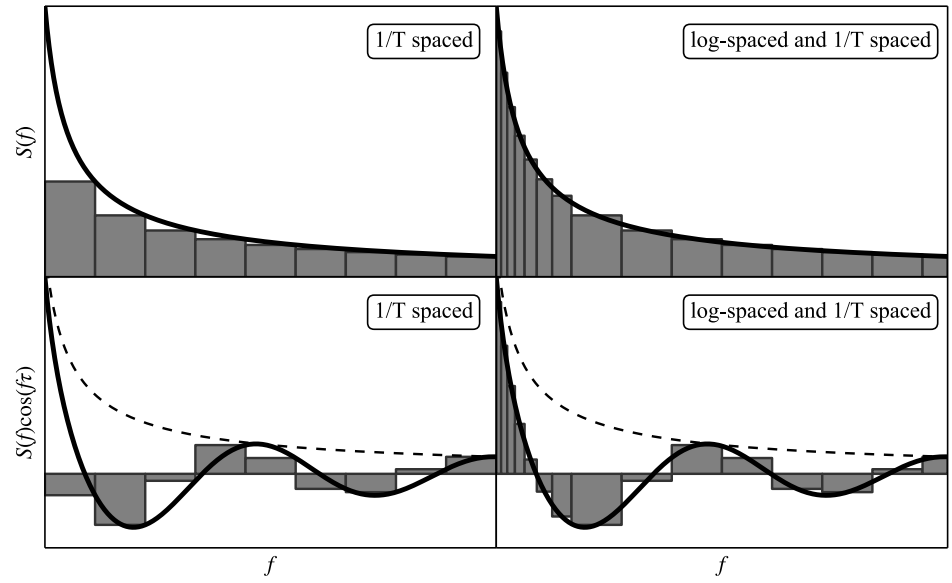

FIG. 3. Approximating the Wiener–Khinchin integral $C(\tau)$ as a sum of $S(f_i)$ over the frequencies f_i, for $\tau = 0$ (top) and for a representative nonzero τ (bottom). For "red" spectra that blow up at low frequencies, a sum with equally spaced bins (left panels) is improved by inserting more bins at low frequencies (right panels).

of τ, by way of Eq. (5) and of our improved scheme.

Specifically, we rewrite the Wiener–Khinchin integral as the sum

$$C(\tau) = \int_0^\infty \mathrm{d}f\, S(f)\cos(2\pi f\tau) \approx \int_{\log f_L}^{\log f_M} \mathrm{d}(\log f)\, fS(f)\cos(2\pi f\tau) + \int_{f_M}^{f_H} \mathrm{d}f\, S(f)\cos(2\pi f\tau) \tag{6}$$

where the low- and high-frequency cutoffs f_L and f_H, as well as the turnover frequency f_M, are chosen depending on the specific spectrum $S(f)$ (or family of spectra) under examination. We approximate the two subintegrals in Eq. (6) with the extended Simpson rule [17]

$$\int_{x_1}^{x_n} g(x)\,\mathrm{d}x \approx \frac{2\Delta x}{6}\Big(g(x_1) + 4g(x_2) + 2g(x_3) + 4g(x_4)\ldots + 4g(x_{n-1}) + g(x_n)\Big) \quad \text{(with odd } n\text{)}, \tag{7}$$

where x is $\log f$ and f respectively for the two subintegrals, and where g and the edges and spacings of the x_i are chosen accordingly. Incorporating both the logarithmic bin size and the Simpson-rule coefficients in the weights w_i, we can then write

$$C(\tau) \approx \sum_i w_i S(f_i)\cos(2\pi f_i\tau), \tag{8}$$

which summarizes our first improved scheme for approximating the covariance matrix.

As shown in Figs. 1 and 2, the low-frequency-logarithmic expansion improves on the Fourier-sum prescription, especially for larger γ. Just as for Fourier sums, accuracy is better for the post-subtraction covariance matrix than for the unsubtracted $C(\tau)$. To produce these plots, the integration limit f_L of Eq. (6) was matched with the cutoff frequency $f_L = 1/(10T)$ of the exact covariance function.

IV. TIME-DOMAIN COSINE EXPANSION OF COVARIANCE

Our second improved scheme is motivated by the observation that the Wiener–Khinchin integrand is very oscillatory: this is challenging for reduced-rank approximations that represent the integral of Eq. (2) as a sum over a small set of evaluations of the integrand, unless $S(f)$ declines so rapidly with f that only a few cycles of $\cos(2\pi f\tau)$ are relevant.

By contrast, the integrated $C(\tau)$ is usually a smooth function of τ. With this in mind, a useful way to interpret Eq. (5) is to view it as a truncated series expansion for the time-domain function $C(\tau)$, rather than a discrete-sum

approximation for the Fourier-domain integral (2). In particular, if we take $f_k = k/2T$, with $k = 0, \ldots, q$, Eq. (5) corresponds to the standard discrete cosine transform [17] with coefficients given by

$$\rho_j = \frac{2}{p} \sum_{k=1}^{p}{}' C\left(\frac{(k-1/2)T}{p}\right) \cos\left(\frac{\pi j(k-1/2)}{p}\right). \tag{9}$$

Here p can be equal to q, or larger, in which case we are using the first $q+1$ terms of the p-point transform. The prime over the summation symbol indicates that ρ_0 must be multiplied by 1/2. The cosine transform can also be viewed as an interpolation of $C(\tau)$ in terms of Chebyshev polynomials $T_j(\xi)$, where $\xi = \cos(\pi\tau/T)$. This interpretation offers some guarantees about the optimality of the cosine-series approximation [17], although these would be expressed relative to the transformed time coordinate ξ.

In fact, we do not even need to work with the integrated $C(\tau)$, but we can obtain the ρ_j directly from the power spectral density. In the limit $p \to \infty$, replacing Eq. (2) in Eq. (9) yields

$$\rho_j = (-1)^{j+1} \int S(f) \frac{4fT\sin(2\pi fT)}{\pi j^2 - 4\pi f^2 T^2} \mathrm{d}f, \tag{10}$$

which describes the coefficients for our second improved scheme, which was first conceived by the second of these authors (MV).

This scheme and the logarithmic-spacing scheme of Sec. III are related to the "sparse spectrum" method of Lázaro–Gredilla et al. [16], in that we seek to optimize the choice of the frequencies in the last row of Eq. (5); however, Lázaro–Gredilla and colleagues do so by straight numerical optimization, while our recipes are motivated by approximating the covariance integral as a finite sum, and by expanding the covariance, seen as a time-domain function, as a truncated cosine series.

The cosine expansion can be used for spectral estimation with a very general representation of $S(f)$. If $S(f)$ is provided as an interpolant $\sum_m S(f_m) I_m(f - f_m)$, where $I_m(f - f_m)$ is an interpolation kernel, such as a boxcar or triangle function for nearest-neighbor and linear interpolation respectively, then the coefficients ρ_j are related to the $S(f_m)$ by the linear transformation

$$\rho_j = M_{jm} S(f_m), \quad \text{with} \quad M_{jm} = (-1)^{j+1} \int I_k(f - f_m) \frac{4fT\sin(2\pi fT)}{\pi j^2 - 4\pi f^2 T^2} \mathrm{d}f, \tag{11}$$

where the matrix M_{jm} is constant once the f_m are chosen.

As shown in Figs. 1 and 2, the cosine expansion is remarkably accurate for the unsubtracted $C(\tau)$, since it includes a constant term and it is not limited to representing periodic functions over $[0, T]$ (rather, it represents functions with periodic first derivative). However, it suffers from oscillatory error (a form of Gibbs phenomenon [18]) that becomes more evident in the post-subtraction covariance. That error is alleviated by replacing T with a small multiple in Eqs. (5) and (10).

V. CONCLUSIONS

In this paper we investigate the low-rank approximation of stationary covariance matrices (Sec. I), with a special focus on applications in pulsar-timing data analysis, and we quantify the accuracy of the commonly used "Fourier-sum" approximation [13] (Sec. II) and of two novel methods proposed here (Secs. III and IV), by way of the matrix norm of the covariance-matrix error. We have included code demonstrations in the Python and C programming languages in the ancillary files of this manuscript.

The Fourier-sum approximation is accurate only when the covariance matrix is projected onto a subspace orthogonal to second-order polynomials in time (which is usually the case in pulsar timing because the linear and quadratic evolution of the pulsar frequency are fit as model parameters), and if the power spectral density represented by the matrix is not too step at low frequencies (for a power law, this corresponds to spectral index $\gamma \lesssim 7$). Even then, the lack of low-frequency components in the Fourier-sum expansion biases the estimation of timing-model parameters such as quadratic spindown; conversely, it requires the artificial addition of second-order polynomials to the timing model when these are not naturally present, as in the case of dispersion-measure variations [19].

We resolve these maladies in two novel schemes. In the first, we generalize the Fourier-sum expansion by interpreting it as the discrete approximation of the covariance integral, which prompts the addition of low-frequency terms with logarithmic spacing. Only a few terms are needed to improve low-frequency coverage. In the second, we repurpose and modify the Fourier-sum expansion to encode the cosine expansion of the time-domain covariance function (or equivalently, as its Chebyshev interpolation in a transformed time coordinate).

Both schemes vastly outperform the Fourier-sum method in approximating the unprojected covariance, as well as the projected covariance for steep red spectra. Which of the two new methods performs best and how many frequency components are needed depends on the specifics of the random process that is being approximated, and should be determined case by case. Compared to the performing matrix algebra on full-sized covariance matrix, our low-rank approximations greatly reduce memory requirements and computational costs, often by orders of magnitude.

Acknowledgments. RvH was supported by NASA Einstein Fellowship grant PF3-140116. MV was supported by the Jet Propulsion Laboratory RTD program. This work was carried out at the Jet Propulsion Laboratory, California Institute of Technology, under contract to the National Aeronautics and Space Administration. Copyright 2014 California Institute of Technology. Government sponsorship acknowledged.

[1] R. M. Shannon and J. M. Cordes, ApJ **725**, 1607 (2010), arXiv:1010.4794 [astro-ph.SR].
[2] R. van Haasteren, Y. Levin, P. McDonald, and T. Lu, Monthly Notices of the Royal Astronomical Society **395**, 1005 (2009), arXiv:0809.0791.
[3] Z. Arzoumanian, A. Brazier, S. Burke-Spolaor, S. J. Chamberlin, S. Chatterjee, J. M. Cordes, P. B. Demorest, X. Deng, T. Dolch, J. A. Ellis, R. D. Ferdman, N. Garver-Daniels, F. Jenet, G. Jones, V. M. Kaspi, M. Koop, M. Lam, T. J. W. Lazio, A. N. Lommen, D. R. Lorimer, J. Luo, R. S. Lynch, D. R. Madison, M. McLaughlin, S. T. McWilliams, D. J. Nice, N. Palliyaguru, T. T. Pennucci, S. M. Ransom, A. Sesana, X. Siemens, I. H. Stairs, D. R. Stinebring, K. Stovall, J. Swiggum, M. Vallisneri, R. van Haasteren, Y. Wang, and W. W. Zhu, ArXiv e-prints (2014), arXiv:1404.1267.
[4] R. van Haasteren and M. Vallisneri, ArXiv e-prints (2014), arXiv:1407.1838 [gr-qc].
[5] W. Coles, G. Hobbs, D. J. Champion, R. N. Manchester, and J. P. W. Verbiest, MNRAS **418**, 561 (2011), arXiv:1107.5366 [astro-ph.IM].
[6] W. Hager, SIAM Review **31**, 221 (1989).
[7] L. N. Trefethen and D. Bau, *Numerical Linear Algebra* (SIAM, 1997).
[8] D. Goldfarb and K. Scheinberg, Mathematical Programming **99**, 1 (2004).
[9] D. Goldfarb and K. Scheinberg, Mathematical Programming **103**, 135 (2005).
[10] A. Smola and S. Vishwanathan, "LDL factorization for rank-k modifications of diagonal matrices," `http://www.stat.purdue.edu/~vishy/papers/SmoVis03a.pdf` (2004), unpublished manuscript.
[11] C. Rasmussen and C. Williams, *Gaussian processes for machine learning* (MIT Press, 2006).
[12] R. van Haasteren and Y. Levin, Monthly Notices of the Royal Astronomical Society **428**, 1147 (2013), arXiv:1202.5932 [astro-ph.IM].
[13] L. Lentati, P. Alexander, M. P. Hobson, S. Taylor, J. Gair, S. T. Balan, and R. van Haasteren, Phys. Rev. D **87**, 104021 (2013), arXiv:1210.3578 [astro-ph.IM].
[14] C. J. Paciorek, Journal of Statistical Software **19**, 1 (2007).
[15] K. Chalupka, C. K. I. Williams, and I. Murray, Journal of Machine Learning Research **14**, 333 (2013).
[16] M. Lázaro-Gredilla, J. Q. Candela, C. E. Rasmussen, and A. R. Figueiras-Vidal, Journal of Machine Learning Research **11**, 1865 (2010).
[17] W. H. Press, S. A. Teukolsky, W. T. Vetterling, and B. P. Flannery, *Numerical recipes in C. The art of scientific computing* (Cambridge University Press, Cambridge, 1992).
[18] R. N. Bracewell, *The Fourier Transform and its Applications*, 2nd ed. (McGraw-Hill, New York, 1978).
[19] L. Lentati, P. Alexander, M. P. Hobson, F. Feroz, R. van Haasteren, K. J. Lee, and R. M. Shannon, MNRAS **437**, 3004 (2014), arXiv:1310.2120 [astro-ph.IM].